arXiv:1401.1693v1 [cs.LO] 8 Jan 2014

Certification for μ calculus with winning strategies

(Presented at the VeriSure workshop associated with CAV'13 in St. Petersburg in July 2013)

Martin Hofmann[1] and Harald Rueß[2]

[1] Department of Informatics, Ludwig-Maximilians-Universität, München, Germany
[2] fortiss An-Institut Technische Universität München, D-80805 München, Guerickestr. 25, Germany

Abstract. We define memory-efficient certificates for μ-calculus model checking problems based on the well-known correspondence of the μ-calculus model checking with winning certain parity games. Winning strategies can independently checked, in low polynomial time, by observing that there is no reachable strongly connected component in the graph of the parity game whose largest priority is odd. Winning strategies are computed by fixpoint iteration following the naive semantics of μ-calculus. We instrument the usual fixpoint iteration of μ-calculus model checking so that it produces evidence in the form of a winning strategy; these winning strategies can be computed in polynomial time in $|S|$ and in space $O(|S|^2|\phi|^2)$, where $|S|$ is the size of the state space and $|\phi|$ the length of the formula ϕ. On the technical level our work can be seen as a new, simpler, and immediate constructive proof of the correspondence between μ-calculus and parity games.

1 Introduction

We develop algorithms for constructing concise certificates for μ-calculus model checking problems and for efficiently and independently checking these certificates by a trustworthy checker.

These developments may form the underpinning for a sound integration of μ-calculus model checking into a verification system such as PVS [18]. Using Shankar's kernel of truth [20] approach, which is based on checking the verification and on verifying the checker, certificates are generated using an untrusted implementation of our μ-calculus model checking algorithms, and certificates are then checked by means of an executable PVS function, which itself is verified in a trusted kernel of PVS. In contrast to logical integration frameworks based on computational reflection (e.g. [3]) the kernel of truth approach does not require proving the correctness of the complete implemention of the verification procedure.

In this way it should be possible to generate checkable certificates for the bisimulation between programs and for model checking problems for both linear time temporal logics and computation tree logics [10] as the basis for assurance cases and certification arguments for safety-critical systems. Moreover, certificates for μ-calculus model checking might also be used a symmetric abstraction-refinement-based model checking engines for the full μ-calculus based on refining over-approximations using *spurious counterexamples* and relaxing under-approximations using *dubious witnesses* along the lines of [23, 22], for sending code together with proofs of arbitrary safety and liveness properties properties, which are then checked by code consumers according to the *proof-carrying code* paradigm of [17], and for synthesizing *correct-by-construction* controllers from these certificates along the lines of [22].

Our main result is an effective and efficient instrumentation of the usual fixpoint iteration of μ-calculus model checking [5] for generating certificates that are independently checkable in low polynomial time. This construction builds on the well-known correspondence of model checking for the μ-calculus with winning certain parity games by Emerson and Jutla [9]. Parity games are equivalent via linear time reductions to the problem of μ calculus model checking (e.g. [11]), Determinacy of parity games follows directly from Martin's most general result on the

determinacy of Borel games [15]. Players of parity games may restrict themselves to memoryless strategies; this also implies that for each vertex one of the players has a winning strategy, so there are no draws. Algorithms for generating witnesses for players of parity games and their complexity are described in [13].

There have been many results and algorithms for constructing witnesses and counterexamples of various forms for different sublogics, including LTL, $ACTL$, CTL, CTL^*, or the μ-calculus [6, 1, 19, 26, 7, 21, 12]. Local model checking procedures for determining whether finite-state systems have properties expressible in the μ-calculus incrementally construct tableau proofs [27, 24, 8], which can be proof-checked independently. The size of the constructed tableaux can be exponential in the number of states of the model. Based on the tableau method of local μ-calculus model checking, Kick [14] proposes an optimized construction by identifying isomorphic subproofs. Namjoshi [16] introduced the notion of *certifying model checker* that can generate independently checkable witnesses for properties verified by a model checker. He defines witnesses for properties of labelled transition systems expressed in the μ-calculus based on parity games over alternating tree automata. His technical developments rely on μ-calculus signatures [25] for termination, and exploits the correspondence between μ-calculus model checking with winning parity games [9].

In contrast to the above methods, the witnesses generated by our global model checking algorithm are rather small, as they can be represented in space in $O(|S|^2|\phi|^2)$, where $|S|$ is the size of the state space and $|\phi|$ is the length of the formula ϕ. On the technical level our work can be seen as a new, simpler, and immediately constructive proof of the correspondence between μ-calculus and parity games. Winning strategies are computed by fixpoint iteration following the naive semantics of μ-calculus. No complex auxiliary devices such as signatures [25] or alternating automata [9] are needed.

The paper is structured as follows. In Sections 2 and 3 we are reviewing standard developments for the μ-calculus to keep the paper as self-contained as possible. The definition of the nesting depth in Section 2, however, is non-standard and central to the technical developments in this paper. The low polynomial-time checker for certificates in Section 3 is inspired by the standard algorithm for nonemptiness of Street automata. Our constructive proof of the correspondence between μ calculus model checking and winning parity games forms the basis of our main contribution, namely the instrumentation of the global model checking iteration for producing memory-efficient certificates in the form of winning strategies.

2 Syntax and Semantics

We are assuming variables $X \in \mathcal{X}$, propositions $P \in \mathcal{P}$, and actions $a \in \mathcal{A}$.

Definition 1. *The set of μ-calculus formulae ϕ is given by the grammar*

$$\phi \ ::= \ X \mid P \mid \langle a\rangle\phi \mid [a]\phi \mid \phi_1 \wedge \phi_2 \mid \phi_1 \vee \phi_2 \mid \mu X.\,\phi \mid \nu X.\,\phi$$

The set of free variables $FV(\phi) \subseteq \mathcal{X}$, the size $|\phi|$ of a formula, and the substitution $\phi[Z := \psi]$ of formula ψ for any free occurrence $Z \in FV(\phi)$ are defined in the usual way.

The notations $Q \in \{\mu, \nu\}$, $M \in \{[a], \langle a\rangle \mid a \in \mathcal{A}\}$, $* \in \{\wedge, \vee\}$ are used to simplify inductive definitions. We define the nesting depth $nd(QX.\phi)$ of a fixpoint formula as one plus the maximal nesting depth—recursively—of all fixpoint formulas encountered until any free occurrence of X

in ϕ. Formally,

$$\begin{aligned} nd(X,\phi) &= 0, \text{ if } X \notin FV(\phi), \text{otherwise:} \\ nd(X,\ X) &= 0 \\ nd(X,\ \phi_1 * \phi_2) &= max(nd(X,\phi_1),\ nd(X,\phi_2)) \\ nd(X,\ M\phi) &= nd(X,\phi) \\ nd(X,\ QY.\,\phi) &= max(nd(QY.\,\phi), nd(X,\ \phi)) \\ nd(QX,\phi) &= 1 + nd(X,\phi) \\ nd(\phi) &= 0, otherwise. \end{aligned}$$

For example, $nd(QW.Y) = 1$ so $nd(QY.X \wedge QW.Y) = 2$ and $nd(QX.QY.X \wedge QW.Y) = 3$, however $nd(X,\ QY.X \wedge QW.Y) = 1$.

The salient property of the nesting depth is summarised by the following lemma which is easily proved by induction.

Lemma 1. *Let $\phi = QX.\psi_1[Z := QY.\psi_2]$ where $X \in FV(\psi_2)$ and $Z \in FV(\psi_1)$; then $nd(QY.\psi_2) < nd(\phi)$.*

Thus, if we travel down from a fixpoint quantifier to an occurrence of its bound variable then all the fixpoint quantifiers encountered on the way in the abstract syntax tree have strictly smaller nesting depth.

The semantics of μ-calculus formulae is given in terms of labelled transition systems (LTS), consisting of a nonempty set of states S, and a family of total relations $\xrightarrow{a} \ \in S \times S$ for each action $a \in \mathcal{A}$ and, finally, an assignment $T \in S \to 2^{\mathcal{P}}$ which tells for each state s which atomic propositions $P \in \mathcal{P}$ are true in that state. If T is an LTS, we use $|T|$ for its set of states; $\xrightarrow{a}_T$ or simply $\xrightarrow{a}$ for its transition relation and T itself for its interpretation of atomic propositions.

Fix a transition system T and put $S = |T|$. For η is a finite partial function from $\mathcal{X}$ to S with $FV(\phi) \subseteq dom(\eta)$ we define $[\![\phi]\!]\eta \subseteq S$ by

$$\begin{aligned} [\![P]\!]\eta &= \{s \mid P \in T(s)\} \\ [\![X]\!]\eta &= \eta(X) \\ [\![\phi_1 \vee \phi_2]\!]\eta &= [\![\phi_1]\!]\eta \cup [\![\phi_2]\!]\eta \\ [\![\phi_1 \wedge \phi_2]\!]\eta &= [\![\phi_1]\!]\eta \cap [\![\phi_2]\!]\eta \\ [\![\langle a\rangle\phi]\!]\eta &= pre(\xrightarrow{a})([\![\phi]\!]\eta) \\ [\![[a]\phi]\!]\eta &= \widetilde{pre}(\xrightarrow{a})([\![\phi]\!]\eta) \\ [\![\mu X.\phi]\!]\eta &= lfp(U \mapsto [\![\phi]\!]\eta[X := U]) \\ [\![\nu X.\phi]\!]\eta &= gfp(U \mapsto [\![\phi]\!]\eta[X := U]) \end{aligned}$$

The sets $pre(\xrightarrow{a})([\![\phi]\!]\eta)$ and $\widetilde{pre}(\xrightarrow{a})([\![\phi]\!]\eta)$ respectively denote the *preimage* and the *weakest precondition* of the set $[\![\phi]\!]\eta$ with respect to the binary relation $\xrightarrow{a}$; formally:

$$\begin{aligned} &s \in pre(\xrightarrow{a})([\![\phi]\!]\eta) \text{ iff } \exists t \in S.\ s \xrightarrow{a} t \text{ and } t \in [\![\phi]\!]\eta \\ &s \in \widetilde{pre}(\xrightarrow{a})([\![\phi]\!]\eta) \text{ iff } \forall t \in S.\ s \xrightarrow{a} t \text{ implies } t \in [\![\phi]\!]\eta \end{aligned}$$

Given the functional $F(U) = [\![\phi]\!]\eta[X := U]$, $lfp(F)$ and $gfp(F)$respectively denote the least and the greatest, with respect to the subset ordering on 2^S, fixpoints of F. These fixpoints exist, since F is monotone.

$$
\begin{aligned}
sem(X,\eta) &= \eta(X)\\
sem(\mu X.\phi,\eta) &= iter_X(\phi,\eta,\emptyset)\\
sem(\nu X.\phi,\eta) &= iter_X(\phi,\eta,S)\\
sem(\phi_1 \wedge \phi_2,\eta) &= sem(\phi_1,\eta) \cap sem(\phi_2,\eta)\\
sem(\phi_1 \vee \phi_2,\eta) &= sem(\phi_1,\eta) \cup sem(\phi_2,\eta)\\
sem([a]\phi,\eta) &= \widetilde{pre}(\xrightarrow{a})(sem(\phi,\eta))\\
sem(\langle a\rangle\phi,\eta) &= pre(\xrightarrow{a})(sem(\phi,\eta))
\end{aligned}
$$

$$
\begin{aligned}
iter_X(\phi,\eta,U) = {} & \mathit{if}\ U = U_p\ \mathit{then}\ U\ \mathit{else}\ iter_X(\phi,\eta,U_p)\\
& \mathit{where}\ U_p := sem(\phi,\eta[X := U])
\end{aligned}
$$

Fig. 1. Fixpoint iteration for computing the semantics of μ-calculus formulas.

Proposition 1. $[\![QX.\phi]\!]\eta = [\![\phi[X := QX.\phi]]\!]\eta$.

For the monotonicity of F, $\emptyset \subseteq F(\emptyset) \subseteq F^2(\emptyset) \subseteq \ldots$ and $S \supseteq F(S) \supseteq F^2(S) \supseteq \ldots$. Moreover, if S is finite then we have

$$
\begin{aligned}
[\![\mu X.\phi]\!]\eta &= \{s \in S \mid \mathit{exists}\ t \leq |S|.\ s \in F^t(\emptyset)\},\\
[\![\nu X.\phi]\!]\eta &= \{s \in S \mid \mathit{forall}\ t \leq |S|.\ s \in F^t(S)\}.
\end{aligned}
$$

Therefore, in the case S is finite, the iterative algorithm in Figure 1 computes $[\![\phi]\!]\eta$.

Proposition 2. $[\![\phi]\!]\eta = sem(\phi,\eta)$.

Lemma 2. *$s \notin [\![\phi]\!]\eta$ iff $s \in [\![\phi^*]\!]\eta'$, where $\eta'(X) = S\backslash\eta(X)$ and ϕ^* is the dual of ϕ given by*

$$
\begin{aligned}
(P)^* &= P\\
(X)^* &= X\\
(\phi_1 \wedge \phi_2)^* &= \phi_1^* \vee \phi_2^*\\
(\phi_1 \vee \phi_2)^* &= \phi_1^* \wedge \phi_2^*\\
([a]\phi)^* &= \langle a\rangle\phi^*\\
(\langle a\rangle\phi)^* &= [a]\phi^*\\
(\mu X.\phi)^* &= \nu X.\phi^*\\
(\nu X.\phi)^* &= \mu X.\phi^*
\end{aligned}
$$

3 Parity Games

A *parity game* is given by the following data:

- a (finite or infinite) set of positions Pos partitioned into proponent's (Player 0) and opponent's (Player 1) positions: $Pos = Pos_0 + Pos_1$;
- a total edge relation $\rightarrow \subseteq\ Pos \times Pos$;[1]

[1] total means forall $p \in Pos$ there exists $p' \in Pos$ with $p \rightarrow p'$.

– a function $\Omega \in Pos \to \mathbb{N}$ with a finite range; we call $\Omega(p)$ the priority of position p.

The players move a token along the edge relation $\to$. When the token is on a position in Pos_0 then proponent decides where to move next and likewise for opponent.

In order to formalize the notion of "to decide" we must introduce strategies. Formally, a strategy for a player $i \in 0, 1$ is a function σ that for any nonempty string $\vec{p} = p(0) \dots p(n)$ over Pos and such that $p(k) \to p(k+1)$ for $k = 0 \dots n-1$ and $p(n) \in Pos_i$ associates a position $\sigma(\vec{p}) \in Pos$ such that $p(n) \to \sigma(\vec{p})$.

Given a starting position p and strategies σ_0 and σ_1 for the two players one then obtains an infinite sequence of positions (a "play") $p(0), p(1), p(2), \dots$ by

$$\begin{aligned} p(0) &= p \\ p(n+1) &= \sigma_i(p(0) \dots p(n)) \ \ \textit{where } p(n) \in Pos_i \end{aligned}$$

We denote this sequence by $play(p, \sigma_0, \sigma_1)$.

The play is won by proponent (Player 0) if the largest number that occurs infinitely often in the sequence $\Omega(play(p, \sigma_0, \sigma_1))$ is even and it is won by opponent if that number is odd. Note that $\Omega(_)$ is applied component-wise and that a largest priority indeed exists since Ω has finite range.

Player i wins from position p if there exists a strategy σ_i for player i such that for any strategy σ_{1-i} for the other player (Player $1-i$) player i wins $play(p, \sigma_0, \sigma_1)$. We write W_i for the set of positions from which Player i wins.

A strategy σ is *positional* if $\sigma(p(0)..p(n))$ only depends on $p(n)$. Player i *wins positionally* from p when the above strategy σ_i can be chosen to be positional.

The following is a standard result.

Theorem 1. *Every position p is either in W_0 or in W_1 and player i wins positionally from every position in W_i.*

In view of this theorem we can now confine attention to positional strategies. A strategy that wins against all positional strategies is indeed a winning strategy (against all strategies) since the optimal counterstrategy is itself positional.

Example 1. Fig. 2 contains a graphical display of a parity game. Positions in Pos_0 and Pos_1 are represented as circles and boxes, respectively, and labelled with their priorities. Formally, $Pos = \{a, b, c, d, e, f, g, h, i\}$ and $Pos_0 = \{b, d, f, h\}$ and $Pos_1 = \{a, c, e, g, i\}$ and $\Omega(a) = 3$, $\dots$, and $\to = \{(a, b), (b, f), \dots\}$.

In the right half of Fig. 2 the winning sets are indicated and corresponding positional winning strategies are given as fat arrows. The moves from positions that are not in the respective winning set are omitted but can of course be filled in in an arbitrary fashion.

3.1 Certification of winning strategies

Given a parity game with finitely many positions, presented explicitly as a finite labelled graph, and a partition of Pos into V_0 and V_1 we are now looking for an easy to verify certificate as to the fact that $V_0 = W_0$ and $V_1 = W_1$.

In essence, such a certificate will consist of a positional strategy σ_i for each player i such that i wins using σ_i from every position p in V_i. Clearly this implies $V_i = W_i$ and the above theorem asserts that in principle such certificates always exist when $V_i = W_i$. However, it remains to explain how we can check that a given positional strategy σ_i wins from a given position p.

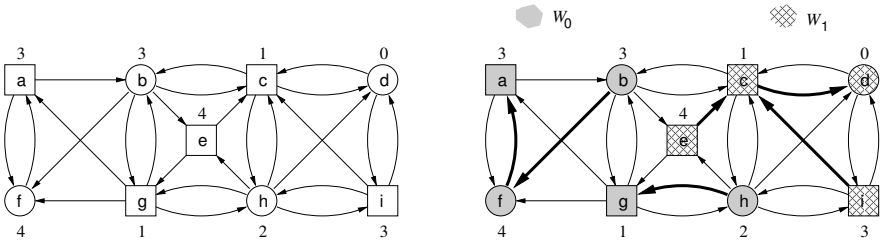

Fig. 2. A parity game and its decomposition into winning sets.

We first note that for this it is enough that it wins against any adversarial positional strategy because the "optimal" counterstrategy, i.e., the one that wins from all adversarial winning positions is positional (by theorem 1). Thus, given a positional strategy σ_i for player i we can remove all edges from positions $p' \in Pos_i$ that are not chosen by the strategy and in the remaining game graph look for a cycle whose largest priority has parity $1 - i$ and is reachable from p. If there is such a cycle then the strategy was not good and otherwise it is indeed a winning strategy for Player i.

Algorithmically, the absence of such a cycle can be checked by starting a depth-first search from every position of adversary priority (parity $1 - i$) after removing all positions of favourable and higher priority (parity i). More efficiently, one can decompose the reachable (from the purported winning set) part of the remaining graph into nontrivial strongly connected components (SCC). If such an SCC only contains positions whose priority has parity $1 - i$ then, clearly, the strategy is bad. Otherwise, one may remove the positions with the largest priority of parity i, decompose the remaining graph into SCCs and continue recursively. Essentially, this is the standard algorithm for nonemptiness of Strett automata described in [2]. This paper also describes an efficient algorithm for SCC decomposition that could be used here.

Example 2. After removing the edges not taken by Player 0 according to his purported winning strategy we obtain the following graph:

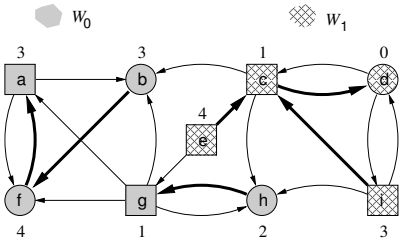

We see that the two reachable SCC from W_0 are $\{a, b, f\}$ and $\{g, h\}$. The first one contains the cycles a, f and a, b, f which both have largest priority 4. The other one is itself a cycle with largest priority 2.

Likewise, adopting the viewpoint of Player 1, after removing the edges not taken by his strategy we obtain

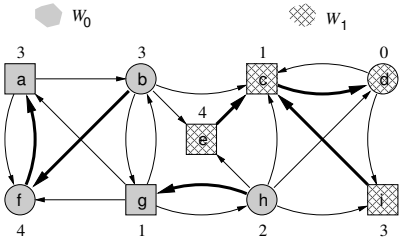

and find the reachable (from W_1) SCCs to be $\{c, d, i\}$. The only cycles therein are d, e and d, e, i. Both are good for Player 1.

3.2 Game-theoretic characterization

The game $G(T, \eta)$ associated with the LTS T and η as above is the defined as follows. Positions are pairs (s, ϕ) where $FV(\phi) \subseteq dom(\eta)$ and $s \in S$. In positions of the form (s, ψ) where ψ starts with $\vee$ or $\langle a \rangle$, it is proponent's (Player 0) turn. The possible moves (for proponent to choose from) are:

$$
\begin{aligned}
(s, \psi_1 \vee \psi_2) &\rightsquigarrow (s, \psi_1) \\
(s, \psi_1 \vee \psi_2) &\rightsquigarrow (s, \psi_2) \\
(s, \langle a \rangle \psi) &\rightsquigarrow (t, \psi) \ \textit{where} \ (s \xrightarrow{a} t) \in T.
\end{aligned}
$$

In positions of the form (s, ψ) where ψ starts with $\wedge$ or $[a]$ it is the opponent's turn. The possible moves (for opponent to choose from) are:

$$
\begin{aligned}
(s, \psi_1 \wedge \psi_2) &\rightsquigarrow (s, \psi_1) \\
(s, \psi_1 \wedge \psi_2) &\rightsquigarrow (s, \psi_2) \\
(s, [a] \psi) &\rightsquigarrow (t, \psi) \ \textit{where} \ (s \xrightarrow{a} t) \in T.
\end{aligned}
$$

In positions of the form $(s, QX.\phi)$ the proponent is to move, but there is only one possible move:

$$
\begin{aligned}
(s, \mu X.\phi) &\rightsquigarrow (s, \phi[X := \mu X.\phi]) \\
(s, \nu X.\phi) &\rightsquigarrow (s, \phi[X := \nu X.\phi])
\end{aligned}
$$

In positions of the form (s, X), (s, P), the proponent is to move, but there is only one possible move: (s, X) (respectively (s, P)) itself. So, de facto, the game ends in such a position.

Each position (s, ϕ) is assigned a natural number $\Omega(s, \phi)$, its priority, as follows:

$$
\begin{aligned}
\Omega(s, \mu X.\phi) &= 2 * nd(\mu X.\phi) + 1 \\
\Omega(s, \nu X.\phi) &= 2 * nd(\nu X.\phi) \\
\Omega(s, P) &= 0 \ \ \textit{if} \ P \in T(s) \\
\Omega(s, P) &= 1 \ \ \textit{if} \ P \notin T(s) \\
\Omega(s, X) &= 0 \ \ \textit{if} \ s \in \eta(X) \\
\Omega(s, X) &= 1 \ \ \textit{if} \ s \notin \eta(X) \\
\Omega(s, \phi) &= 0 \ \ \textit{in all other cases}.
\end{aligned}
$$

For any position (s, ϕ) we can consider the subgame consisting of the positions reachable from (s, ϕ). Even if T is infinite this subgame has only finitely many priorities because the only second components occurring in reachable positions are subformulas of ϕ and one-step unwindings thereof. This subgame is therefore a parity game to which the previous section applies.

Example 3. Let $\phi = \mu X.\, P \vee \langle a \rangle X$ which asserts that a state where P is true can be reached.

Define the transition system T by $|T| = \{s, t\}$ and $\xrightarrow{a}_T = \{(s,s), (s,t), (t,t)\}$ and $T(s) = \emptyset$ and $T(t) = \{P\}$. The associated game graph is as follows

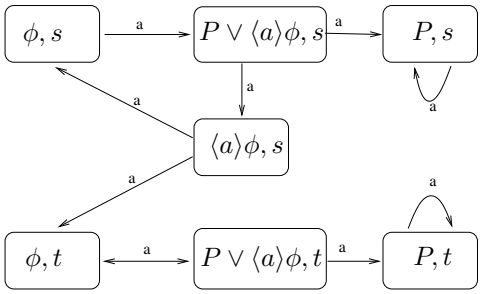

The priorities of the positions labelled $(\phi, s), (\phi, t), (P, s)$ are 1; the priorities of the three other positions are 0.

Player 0 wins from every position except (P, s). The winning strategy moves to $(\langle a \rangle \phi, s)$ and then (ϕ, t) and then (P, t). Note that a strategy that moves from $(P \vee \langle a \rangle \phi, s)$ to (ϕ, s) looses even though it never leaves the winning set W_0. Thus, in order to compute winning strategies it is not enough to choose any move that remains in the winning set.

Theorem 2. *If $s \in [\![\phi]\!]\eta$ then proponent wins $G(T, \eta)$ from (s, ϕ).*

Before proving this, we note that the converse is in this case actually a relatively simple consequence.

Corollary 1. *If proponent wins $G(T, \eta)$ from (s, ϕ) then $s \in [\![\phi]\!]\eta$.*

Proof. Suppose that proponent wins $G(T, \eta)$ from (s, ϕ) and $s \notin [\![\phi]\!]\eta$. We then have $s \in [\![\phi^*]\!]\eta'$ using Lemma 2 for the formal dualisation for formulas and complementation for environments. Thus, by the theorem, proponent wins $G(T, \eta')$ from (s, ϕ^*). However, it is easy to see that a winning strategy for proponent in $G(T, \eta')$ from (s, ϕ^*) is tantamount to a winning strategy for opponent in $G(T, \eta)$ from (s, ϕ); so we get a contradiction using theorem 1.

Proof (of Theorem 2). The proof of Theorem 2 now works by induction on ϕ. The cases where ϕ is a formula with an outermost fixpoint are the interesting ones.

In case ϕ is of the form $\mu X.\psi$, define

$$U := \{t \mid \textit{proponent wins } G(T, \eta) \textit{ from } (t, \mu X.\psi)\}.$$

We must show that $[\![\phi]\!]\eta \subseteq U$. By definition of $[\![\phi]\!]\eta$ it suffices to show that $[\![\psi]\!]\eta[X \mapsto U] \subseteq U$. Thus, suppose that $t \in [\![\psi]\!]\eta[X \mapsto U]$. By the induction hypothesis this means that proponent wins $G(T, \eta[X \mapsto U])$ from (t, ψ). Call the corresponding winning strategy σ. We should prove that proponent also wins from $(t, \mu X.\psi)$. We move to $(t, \psi[X := \mu X.\psi])$ and play according to σ until we reach a state $(t', \mu X.\psi)$ at which point we know that $t' \in U$ so we can then continue to play according to the strategy embodied in that statement. Of course, if we never reach such a position then σ will win the whole game.

In case ϕ is of the form $\nu X.\psi$ define $U := [\![\nu X.\psi]\!]\eta$. We define a winning strategy for positions of the form $(t, \nu X.\psi)$ where $t \in U$ as follows. First, we move (forcedly) to $(t, \psi[X := \nu X.\psi])$. We know that $t \in [\![\psi]\!]\eta[X \mapsto U]$ by unwinding so that, inductively, we have a strategy that allows us to either win rightaway, or move to another position $(t', \nu X.\psi)$ where $t' \in U$ and all priorities encountered on the way are smaller than the one of $\nu X.\psi$ due to the definition

of nesting depth and in particular Lemma 1. We start over and unless we eventually do win rightaway at some point we would have seen the priority of $\nu X.\psi$ infinitely often which is the largest and even.

We remark that while the previous result is well-known the proof presented here is quite different from the ones in the standard literature, e.g. [4] which uses the order-theoretic concept of signatures and are less compositional than ours in the sense that the proof is not directly by structural induction on formulas but rather on the global development of all the fixpoints.

4 Computing winning strategies via fixpoint iteration

4.1 Fixpoint iteration

It is well-known that the fixpoint iteration in Figure 1 computes $[\![\phi]\!]\eta$ in the finite case. Our goal is to somehow instrument this algorithm so that it produces evidence in the form of a winning strategy. In instrumenting this algorithm to produce evidence in the form of a winning strategy it is not enough to simply compute the winning sets using $sem(_,_)$ and then simply choose moves that do not leave the winning set. This is because of examples like 3 which show that a strategy that never leaves the winning set may nonetheless be losing.

Instead we will use the construction from the proof of Theorem 2. Some care needs to be taken with the exact setup of the typing; in particular, our algorithm will return partial winning strategies (that win on a subset of the whole winning set) but only require sets of states (rather than partial winning strategies) as the values of free variables.

4.2 Computing winning strategies

Partial winning strategies. A *partial winning strategy* is a partial function Σ mapping positions of the game $G(T,\eta)$ to elements of S extended with $\{1,2,*\}$; it must satisfy the following conditions:

STAR If $\Sigma(\phi,s) = *$ then all immediate successors of (ϕ,s) are in $dom(\Sigma)$;
OR If $\Sigma(\phi,s) = i \in \{1,2\}$ then ϕ is of the form $\phi_1 \vee \phi_2$ and $(\phi_i,s) \in dom(\Sigma)$;
DIA If $\Sigma(\phi,s) = s' \in S$ then ϕ is of the form $\langle a\rangle\psi$ and $s \xrightarrow{a} s'$ and $(\psi,s') \in dom(\Sigma)$.
WIN Player 0 wins from all the positions in $dom(\Sigma)$ and the obvious strategy induced[2] by Σ is a winning strategy for Player 0 from those positions.

Note that the empty function is in particular a partial winning strategy. To illustrate the notation we describe a (partial) winning strategy for the entire winning set for Example 3:

$$\begin{aligned}
\Sigma(\phi,s) &= * \\
\Sigma(P \vee \langle a\rangle\phi,s) &= 2 \\
\Sigma(\langle a\rangle\phi,s) &= t \\
\Sigma(\phi,t) &= * \\
\Sigma(P \vee \langle a\rangle\phi,t) &= 1 \\
\Sigma(P,t) &= *, \ \textit{and undefined elsewhere.}
\end{aligned}$$

So, $dom(\Sigma) = \{(\phi,s),\ldots,(P,t)\}$ and, indeed, player 0 wins from all these positions by following the advice given by Σ. Of course, $\Sigma'(P,t) = *$ and undefined elsewhere is also a partial winning strategy albeit with smaller domain of definition.

[2] Arbitrary moves outside $dom(\Sigma)$ + remove the *-setting.

Updating of winning strategies. Suppose that Σ and Σ' are partial winning strategies. A new partial winning strategy $\Sigma + \Sigma'$ with $dom(\Sigma + \Sigma')$ is defined by

$$(\Sigma + \Sigma')(\phi, s) = \mathit{if}\ (\phi, s) \in dom(\Sigma)\ \mathit{then}\ \Sigma(\phi, s)\ \mathit{else}\ \Sigma'(\phi, s).$$

Lemma 3. *$\Sigma + \Sigma'$ is a partial winning strategy and $dom(\Sigma + \Sigma') = dom(\Sigma) \cup dom(\Sigma')$*

Proof. A play following $\Sigma + \Sigma'$ will eventually remain in one of Σ or Σ'; this, together with the fact that initial segments do not affect the outcome of a game implies the claim.

Let ϕ be a formula and let Σ be a partial winning strategy for $G(T, \eta[X \mapsto S])$ such that the mapping $(\rho, s) \mapsto (\rho[X := \psi], s)$ is injective on $dom(\Sigma)$. Furthermore, let Σ' be a winning strategy for $G(T, \eta)$ such that $(\phi, s) \,|\, s \in S$ is contained in $dom(\Sigma')$. A new partial strategy $\Sigma[X := \phi, \Sigma']$ is defined by

$$\begin{aligned}\Sigma[X := \phi, \Sigma'](\rho, s) = \mathit{if}\ (\rho, s) \in dom(\Sigma')\ \mathit{then}\ \Sigma'(\rho, s)\ \mathit{else}\\ \mathit{if}\ \mathit{exists}\ \psi, \rho = \psi[X := \phi]\ \mathit{then}\ \Sigma(\psi, s)\ \mathit{else}\ \mathit{undef}\end{aligned}$$

Lemma 4. *Under the assumptions made the strategy $\Sigma[X := \phi, \Sigma']$ is indeed a partial winning strategy for the game $G(T, \eta)$ and $\{(\rho[X := \phi], s) | (\rho, s) \in dom(\Sigma)\} \subseteq dom(\Sigma[X := \phi, \Sigma'])$.*

Proof. Injectivity of the substitution $[X := \phi]$ shows that $\Sigma[X := \phi, \Sigma']$ is well defined. A game according to $\Sigma[X := \Sigma']$ starting from $\rho[X := \phi]$ either stays completely in Σ or else reaches one of the positions (X, s) at which point Σ' takes over.

4.3 Computing winning strategies by fixpoint iteration

For any formula ϕ and environment η with $dom(\eta) \supseteq FV(\phi)$ we define a partial winning strategy $\mathit{SEM}(\phi)_\eta$ by the following clauses:

$$\begin{aligned}\mathit{SEM}(X)_\eta &= \lambda\rho, s.\ \mathit{if}\ \rho = X\ \mathit{and}\ s \in \eta(X)\ \mathit{then}\ *\ \mathit{else}\ \mathit{undef}\\ \mathit{SEM}(P)_\eta &= \lambda\rho, s.\ \mathit{if}\ \rho = X\ \mathit{and}\ P \in T(s)\ \mathit{then}\ *\ \mathit{else}\ \mathit{undef}\end{aligned}$$

$$\begin{aligned}\mathit{SEM}(\phi \wedge \psi)_\eta &= \mathit{SEM}(\phi)_\eta\\ &+ \mathit{SEM}(\psi)_\eta\\ &+ \lambda\rho, s.\ \mathit{if}\ \rho = \phi \wedge \psi\ \mathit{and}\ (s, \phi) \in dom(\mathit{SEM}(\phi)_\eta)\ \mathit{and}\ (s, \psi) \in dom(\mathit{SEM}(\phi)_\eta)\\ &\qquad \mathit{then}\ *\ \mathit{else}\ \mathit{undef}\end{aligned}$$

$$\begin{aligned}\mathit{SEM}(\phi \vee \psi)_\eta &= \mathit{SEM}(\phi)_\eta\\ &+ \mathit{SEM}(\psi)_\eta\\ &+ \{(\phi \vee \psi, s) \mapsto 1 \,|\, (\phi, s) \in dom(\mathit{SEM}(\phi)_\eta)\}\\ &+ \{(\phi \vee \psi, s) \mapsto 2 \,|\, (\psi, s) \in dom(\mathit{SEM}(\phi)_\eta)\}\end{aligned}$$

$$\begin{aligned}\mathit{SEM}([a]\phi)_\eta &= \mathit{SEM}([a]\phi)_\eta\\ \mathit{SEM}(\langle a\rangle\phi)_\eta &= \mathit{SEM}(\phi)_\eta\\ &+ \{(\langle a\rangle\phi, s) \mapsto (\phi, s') \,|\, (\phi, t) \in dom(\mathit{SEM}(\phi)_\eta), s \xrightarrow{a} s'\}\end{aligned}$$

$$
\begin{aligned}
\mathit{SEM}(\mu X.\phi)_\eta &= \Sigma_k, \ \textit{where} \\
&\quad \Sigma_0 = 0; \\
&\quad \Sigma_{n+1} = \mathit{SEM}(\phi)_{[\eta[X \mapsto \{s \mid (\phi,s) \in dom(\Sigma_n)\}]]}[X := \Sigma_n] \\
\mathit{SEM}(\nu X.\phi)_\eta &= \Sigma', \ \textit{where} \\
&\quad U = [\![\nu X.\phi]\!]\eta \\
&\quad \Sigma = \mathit{SEM}(\phi)_{\eta[X \mapsto U]} \\
&\quad \textit{and } \Sigma' \textit{ is obtained from } \Sigma \textit{ by} \\
&\quad\quad 1.\ \textit{removing all positions } (X, s) \\
&\quad\quad 2.\ \textit{redirecting all edges leading into } (X, s) \textit{ into } (\nu X.\phi, s) \\
&\quad\quad 3.\ \textit{substituting } X \textit{ by } \nu X.\phi
\end{aligned}
$$

The following Lemma and Theorem are now immediate from these definitions and the Lemmas 3 and 4.

Lemma 5. $\{s \mid (\phi, s) \in dom(\mathit{SEM}(\phi)_\eta)\} \ = \ [\![\phi]\!]\eta$

Theorem 3. *$\mathit{SEM}(\phi)_\eta$ is a winning strategy for $G(\phi, \eta)$.*

Proposition 1 *$\mathit{SEM}(\phi)_\eta$ can be computed in polynomial time in $|S|$ and in space $O(|S|^2 * |\phi|^2)$, where $|S|$ is the size of the state space and $|\phi|$ the length of the formula ϕ.*

Proof. The computation of $\mathit{SEM}(\phi)_\eta$ follows the one of $sem(\phi, \eta)$ hence the time bound. Just like in the usual implementations of fixpoint iteration one only needs to remember the result of the last iteration. Storing a single partial winning strategy requires space $O(|S|^2.|\phi|)$ and the recursion stack has length $O(\phi)$ thus requiring us to store $O(\phi)$ partial winning strategies at any one time. This yields the announced space bound.

References

1. A. Biere, Y. Zhu, and E. Clarke. Multiple state and single state tableaux for combining local and global nodel checking. In E.-R. Olderog and B. Steffen, editors, *Correct System Design*, volume 1710 of *Lecture Notes in Computer Science*, pages 163–179. Springer, 1999.
2. R. Bloem, H. N. Gabow, and F. Somenzi. An algorithm for strongly connected component analysis in *n* log *n* symbolic steps. *Formal Methods in System Design*, 28(1):37–56, 2006.
3. S. Boutin. Using reflection to build efficient and certified decision procedures. In *Theoretical Aspects of Computer Software*, pages 515–529. Springer, 1997.
4. J. Bradfield and C. Stirling. Modal logics and mu-calculi: An introduction, 2001. available at citeseerx.
5. J. Bradfield and C. Stirling. Modal mu-calculi. *Studies in Logic and Practical Reasoning*, 3:721–756, 2007.
6. E. Clarke, O. Grumberg, K. McMillan, and X. Zhao. Efficient generation of counterexamples and witnesses in symbolic model checking. In *Proceedings of the 32nd annual ACM/IEEE Design Automation Conference*, pages 427–432. ACM, 1995.
7. E. Clarke, S. Jha, Y. Lu, and H. Veith. Tree-like counterexamples in model checking. In *Logic in Computer Science, 2002. Proceedings. 17th Annual IEEE Symposium on*, pages 19–29. IEEE, 2002.
8. R. Cleaveland. Tableau-based model checking in the propositional mu-calculus. *Acta Informatica*, 27(8):725–747, 1990.
9. E. Emerson and C. Jutla. Tree automata, mu-calculus and determinacy. In *Proceedings of the 32nd Annual Symposium on Foundations of Computer Science (FOCS'91)*, pages 368–377. IEEE, 1991.

10. E. Emerson, C. Jutla, and A. Sistla. On model-checking for fragments of -calculus. In C. Courcoubetis, editor, *Computer Aided Verification*, volume 697 of *Lecture Notes in Computer Science*, pages 385–396. Springer Berlin / Heidelberg, 1993.
11. E. Grädel. Back and forth between logic and games. In K. Apt and E. Grädel, editors, *Lectures in Game Theory for Computer Scientists*, pages 99–138. Cambridge University Press, 2011.
12. A. Gurfinkel and M. Chechik. Proof-like counter-examples. In *Proceedings of the 9th international conference on Tools and algorithms for the construction and analysis of systems*, pages 160–175. Springer-Verlag, 2003.
13. M. Jurdzinski. Algorithms for solving parity games. In K. Apt and E. Grädel, editors, *Lectures in Game Theory for Computer Scientists*, pages 74–98. Cambridge University Press, 2011.
14. A. Kick. Generation of counterexamples for the μ-calculus. Technical Report ira-tr-1995-37, Universität Karlsruhe, Germany, 1995.
15. D. A. Martin. Borel determinacy. *The annals of Mathematics*, 102(2):363–371, 1975.
16. K. Namjoshi. Certifying model checkers. In *Computer Aided Verification*, pages 2–13. Springer, 2001.
17. G. Necula. Proof-carrying code. In *Proceedings of the 24th ACM SIGPLAN-SIGACT symposium on Principles of programming languages*, pages 106–119. ACM, 1997.
18. S. Owre, J. M. Rushby, and N. Shankar. Pvs: A prototype verification system. In *Automated DeductionCADE-11*, pages 748–752. Springer, 1992.
19. D. Peled, A. Pnueli, and L. Zuck. From falsification to verification. *FST TCS 2001: Foundations of Software Technology and Theoretical Computer Science*, pages 292–304, 2001.
20. N. Shankar. Rewriting, inference, and proof. In *Rewriting Logic and Its Applications*, pages 1–14. Springer, 2010.
21. N. Shankar and M. Sorea. Counterexample-driven model checking (revisited version). Technical Report SRI-CSL-03-04, SRI International, 2003.
22. M. Sorea. *Verification of real-time systems through lazy approximations*. PhD thesis, University of Ulm, Germany, 2004.
23. M. Sorea. Dubious witnesses and spurious counterexamples. UK Model Checking Days, York, 2005.
24. C. Stirling and D. Walker. Local model checking in the modal mu-calculus. In *TAPSOFT'89*, pages 369–383. Springer, 1989.
25. R. Streett and E. Emerson. The propositional mu-calculus is elementary. In J. Paredaens, editor, *Automata, Languages and Programming*, volume 172 of *Lecture Notes in Computer Science*, pages 465–472. Springer Berlin / Heidelberg, 1984.
26. L. Tan and R. Cleaveland. Evidence-based model checking. In E. Brinksma and K. Larsen, editors, *Computer Aided Verification*, volume 2404 of *Lecture Notes in Computer Science*, pages 641–680. Springer Berlin / Heidelberg, 2002.
27. G. Winskel. A note on model checking the modal [nu]-calculus. *Theoretical Computer Science*, 83(1):157–167, 1991.